

WE CLAIM:

1. A data storage system comprising:
a plurality of storage nodes, each node existing at
a physical location and having one or more associated
5 contexts
interface mechanisms coupled to each storage node
for communicating storage access requests with the
storage node; and
data storage management processes that select one or
10 more of the storage nodes to serve a data storage request
based at least in part upon the particular contexts
associated with each of the storage nodes.
2. The system of claim 1 wherein the data storage
management processes comprise computer-implemented
processes executing in at least one of the storage nodes.
3. The system of claim 1 wherein the data storage
management processes comprise computer-implemented
processes executing in all of the storage nodes.
4. The system of claim 1 wherein the data storage
requests are associated with a set of desired criteria.
5. The system of claim 1 wherein the data storage
requests are associated with a set of desired criteria
and the data storage management processes comprise
processes for matching the desired criteria to the
5 contexts of the storage nodes.
6. The data storage system of claim 1 wherein the
data storage management processes present a unitary
logical volume of data storage to external devices
generating the storage access requests to the selected
5 one or more storage nodes.

7. The data storage system of claim 6 wherein the selected one or more storage nodes are selected such that the contexts associated with the unitary logical volume satisfies the desired criteria associated with the one or
5 more storage nodes.

8. The data storage system of claim 1 wherein the context comprises a political context.

9. The data storage system of claim 1 wherein the context comprises an economic context.

10. The data storage system of claim 1 wherein the context comprises a geographic context.

11. The data storage system of claim 1 wherein the context comprises a network topological context.

12. The data storage system of claim 1 further comprising encryption mechanisms coupled to the interface mechanisms for encrypting storage access requests during communication between nodes.

13. The data storage system of claim 1 further comprising authentication mechanisms coupled to the interface mechanisms for authenticating storage nodes before communicating storage requests.

14. A method of managing distributed data storage comprising the acts of:

providing a plurality of distributed storage nodes;
receiving a data storage task in one of the storage
5 nodes;
determining desired criteria associated with the received unit of data;

selecting one or more of the plurality of storage nodes having an associated context satisfying the desired
10 criteria; and

executing the storage task in the one or more selected storage nodes.

15. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes where neither of the at least two storage nodes individually satisfy the desired performance characteristics, but
5 collectively the at least two storage nodes satisfy the desired performance characteristics.

16. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different geographical locations.

17. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different areas of a single data center.

18. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are connected via different network backbones in a single data center.

19. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different data centers.

20. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different cities.

21. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different political jurisdictions.

22. The method of claim 14 wherein the selection is based upon socio-economic attributes of the physical location of the data storage node.

23. The method of claim 14 wherein the selecting step further comprises matching the desired performance criteria to a context associated with the one or more storage nodes.

24. The method of claim 14 wherein the act of storing comprises storing the data according to a distributed parity scheme analogous to parity distribution found in RAID subsystems.

25. The method of claim 24 wherein the parity paradigm comprises an N-dimensional parity mechanisms where "N" is greater than three.

26. The method of claim 14 wherein the act of storing comprises storing the data in a manner such that the data stored in any one storage node cannot be used in any meaningful fashion without the availability of some
5 or all of the data stored in other storage nodes.

27. A data storage service comprising:
receiving data storage access requests from a file system;

maintaining a state information data structure
5 including state information describing the contexts of a number of network-accessible storage devices; and

using the state information to allocate capacity within the network-accessible storage devices to handle the received data storage request.

28. The data storage service of claim 27 wherein the act of maintaining the state information data structure comprises:

5 detecting a change in state information associated with at least one of the network-accessible storage devices; and

updating the state information associated with the at least one network-accessible storage device to include the change in state information.

29. The data storage service of claim 27 further comprising:

5 dynamically re-allocating capacity within the network-accessible storage devices in response to detecting a change in their associated contexts.

30 The data storage service of claim 29 wherein the dynamic re-allocating is done in the absence of an externally generated data storage access request.

31. A method of transporting data from a source to a destination comprising:

5 determining a context in which each network device that must handle the data transport exists; and selecting a set of network devices such that the data is compliant with the context of each device that handles the data transport.